A Local-Global LDA Model for Discovering Geographical Topics from Social Media

Siwei Qiang
Shanghai Jiao Tong University
qiangsiwei@sjtu.edu.cn

Yongkun Wang
Shanghai Jiao Tong University
ykw@sjtu.edu.cn

Yaohui Jin
Shanghai Jiao Tong University
jinyh@sjtu.edu.cn

arXiv:1607.05806v1 [cs.IR] 20 Jul 2016

ABSTRACT

Micro-blogging services can track users' geo-locations when users check-in their places or use geo-tagging which implicitly reveals locations. This "geo tracking" can help to find topics triggered by some events in certain regions. However, discovering such topics is very challenging because of the large amount of noisy messages (e.g. daily conversations). This paper proposes a method to model geographical topics, which can filter out irrelevant words by different weights in the local and global contexts. Our method is based on the Latent Dirichlet Allocation (LDA) model but each word is generated from either a local or a global topic distribution by its generation probabilities. We evaluated our model with data collected from Weibo, which is currently the most popular micro-blogging service for Chinese. The evaluation results demonstrate that our method outperforms other baseline methods in several metrics such as model perplexity, two kinds of entropies and KL-divergence of discovered topics.

CCS Concepts

•**Information Storage and Retrieval** → Information filtering; Clustering;

Keywords

Geolocation, Geographical topics; Topic modeling; Latent Dirichlet Allocation

1. INTRODUCTION

Micro-blogging services such as Twitter and Weibo are very popular for people to share their status with geo-locations. These data with geo-locations can be easily collected from the web. Such textual data with geo-coordinates or geo-tagged locations usually contain landmark information (e.g., scenic spots or famous restaurants) or information on local events (e.g., movies, vocal concerts, exhibitions or sports games), and hence, can provide us rich and interpretable semantics on different locations. It is also possible to infer inherent geographic variability of topics across various locations.

There are a lot of nice researches on addressing the questions of how the information is created and shared in different geographic locations and how the spatial and linguistic characteristics of people vary across regions. However, the challenge is that the messages with geo-locations are mixed with overwhelming noisy messages of daily chats or expressions of personal emotions, which have little or no relations to the location context. For example, in the city of Shanghai, the Bund (or *Waitan*) is a famous waterfront and one of the most popular scenic spots for tourists. However, even in such a spot, Weibo are still full of daily conversations and greetings such as 'Good night' or 'Have a nice weekend', which has no local semantics. Therefore, it is very difficult to discover meaningful geo-location topics by existing methods such as inferring occurrences of words from local posts.

This paper proposes an effective method to handle noisy messages and model geographical topics of different locations. The proposed method is based on the Latent Dirichlet Allocation (LDA) [1] topic model. The intuitive idea is, for a specific location, the words used by users are different between a) daily conversations; and b) the description of the landmark or local events. The former is relatively consistent across different locations, and denoted as *global* context, while the words in b) varies by sites and are denoted as *local* context. Our method takes the local and global contexts into consideration and models each word to be generated from either its local or the global topic distribution by its estimated probabilities. The proposed strategy is able to distinguish locally featured words from noises and improve the quality of discovered topics. Our evaluation shows its effectiveness for information filtering and geographical topic discovery.

2. RELATED WORK

Two lines of related research are briefly reviewed. Firstly, in order to support location-aware information retrieval or to compare topics across geographical locations, several researches studies have focused on discovering topics from geographical regions [2–6]. For example, Yin et al. proposes and compares three ways of modeling geographical topics, including a location-driven model, a text-driven model, and a joint model called LGTA [3]. Hong et al. takes the Markovian nature of users' locations into account and identifies topics based on location and language [4]. However, they all focused on finding topically coherent locations only, and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2016 ACM. ISBN .
DOI:

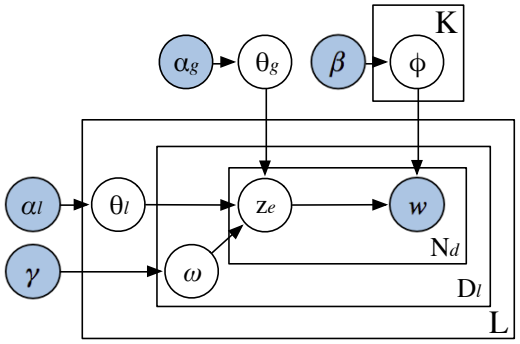


Figure 1: Graphical representation of the proposed Local-Global LDA model.

cannot cope with the difficulty of handling noisy messages.

The second line relevant to our research is local word detection. Based on the premise that local words should have concentrated spatial distributions around their location centers, Backstrom et al. proposes a spatial variation model for analyzing geographic distribution of terms in search engine query logs [7], and this method has been used by Cheng et al. to decide whether a word is local or not [8]. However, it is demonstrated by Wu et al. that this method can be erroneous since it assumes one peak density distribution while many local words can have multiple peaks [9]. Our work is different from the existing works in that, word locality is not generated directly but evaluated by the generation probability, and thus a word (e.g. car) can be both non-*local* for a majority of locations and also *local* for a few particular locations (e.g. automobile 4S shops).

3. METHOD

3.1 Local-Global LDA Model

In our scenario, each geo-located document d is explicitly tagged with a location l, and contains a set of words w_d. A geographical topic z is a meaningful theme shared by similar locations, and each location is associated with a topic distribution $p(z|l)$.

We formalize our model based on the following intuitions. Firstly, words close in space are likely to be clustered into the same geographical topic. Therefore, topics are generated from locations instead of individual documents. Secondly, locally featured words have a more compact geographical scope. For example, 'bravo' is a word for a performer, so that it is more possible to be used at a theater, a concert or a stadium rather than other places. On the contrary, noisy words (e.g., happy, love, city) can have a much wider spatial range. However, some words could be local for certain locations although these words are commonly used at many places. In our method, the role (local or non-local) of a word is determined by its generation probabilities of its local and global semantic contexts. Therefore, our model is named as Local-Global LDA model (or LGLDA).

The graphical representation of our model is shown in Figure 1. Shaded nodes indicate observed variables or priors, while light ones represent latent variables. In order to keep a small set of parameters for simplification, in our model there are one shared set of topics with two different distributions θ_l and θ_g for the local topics and the global topics, respectively. It might be interesting and reasonable to utilize two kinds of ϕ for words' local and global distributions corresponding to the two topic distributions, and we would like to study it in our future work. For a collection of L locations, geo-tagged by D documents, each contains N words, the topic of each word can either be drawn from θ_l or from θ_g. Topic assignment is denoted by z_e, and $e(= l/g)$ indicates whether it is drawn from the local or the global. Since micro-blog is length limited, it is likely to have focused concept. Therefore, if a document is location relevant, each word in it is more likely to be relevant. This relevance of each document is indicated by ω, with binomial distribution prior γ. Finally, the word distribution in K topics is denoted by ϕ. α_l, α_g, β are priors for θ_l, θ_g and ϕ, respectively.

In order to weight between local and global distributions, we add an additional parameter, named local-global weight ratio. Assume the local and global topic distributions are $\theta_l = [p_{l,1}, ..., p_{l,K}]$ and $\theta_g = [p_{g,1}, ..., p_{g,K}]$ respectively, topic assignment is drawn from a concatenated distribution θ in Equation (1).

$$\begin{aligned}\theta &= \frac{\lambda}{\lambda+1}\theta_l p(e=l|w) \oplus \frac{1}{\lambda+1}\theta_g p(e=g|w) \\ &= [\frac{\lambda p_{l,1} p_w^{(l)}}{\lambda+1}, ..., \frac{\lambda p_w^{(l)}}{\lambda+1}, \frac{p_{g,1} p_w^{(g)}}{\lambda+1}, ..., \frac{p_{g,K} p_w^{(g)}}{\lambda+1}]\end{aligned} \tag{1}$$

When λ is too large, the global word set is narrowed and ineffective for noise filtering, while when λ is too small, the size of local words is sparse and it fails to discover meaningful topics. Therefore, an appropriate λ is crucial. In our experiment, it is optimized by estimating the model's perplexity (as illustrated in Figure 2 in Section 4.4).

3.2 Model Inference

Like most Bayesian models, collapsed Gibbs sampling was used for model inference. We present the conditional probability of its latent variables $z_e, \theta_l, \theta_g, \phi$ and w for sampling. Details are omitted for limited space. It is assumed that topic distributions θ_l, θ_g and word distribution ϕ of each topic k are drawn from dirichlet distributions of their respective priors α_l, α_g and β, while locality relevance ω are drawn from a binomial distribution with prior γ.

The conditional probability for sampling the topic assignment z_e of each word is computed in Equation (2), where $z_{e,i} = k$ and $e_i = \kappa$ represent the assignments of the ith word to topic k, and mark the word as local if $\kappa = 1$ otherwise non-local. $\lambda_\kappa = \frac{\lambda}{\lambda+1}(\kappa = l)$ or $\frac{1}{\lambda+1}(\kappa = g)$. $n_{-i,\kappa,k}$ represents the word count with locality assignment κ and topic assignment k, and $-i$ means not including the ith word.

$$\begin{aligned}p(z_{e,i} = k, e_i = \kappa | z_{-i}, e_{-i}) \propto \lambda_\kappa \cdot \frac{n_{-i,\kappa,k}^{(d_i)} + \gamma_\kappa}{n_{-i,\cdot,k}^{(d_i)} + \gamma_l + \gamma_g} \\ \cdot \frac{n_{-i,\kappa,k}^{(l_i)} + \alpha_\kappa}{n_{-i,\kappa,\cdot}^{(l_i)} + K\alpha_\kappa} \cdot \frac{n_{-i,\kappa,k}^{(w_i)} + \beta}{n_{-i,\kappa,k}^{(\cdot)} + W\beta}, \kappa \in \{l, g\}\end{aligned} \tag{2}$$

Consequently, the topic distribution of each location can be computed in Equation (3).

$$p(z_i = k|l_i) = \frac{n_{1,k}^{(l_i)} + \alpha_l}{n_{1,\cdot}^{(l_i)} + K\alpha_l} \quad (3)$$

4. EVALUATION

4.1 Dataset

The effectiveness of the proposed method was evaluated with the data (all in Chinese) collected from Weibo [10]. Without loss of generality, we only focus on messages in Shanghai (the largest city in China and one of the largest cities in the world by population) in 2015. Since most of the locations have been tagged for only a few times, we only selected those locations with a considerable number of posts within a pre-defined spatial range. The original data was preprocessed by filtering out stop words, then nouns and verbs were extracted as valid words with a Chinese POS(part-of-speech) tagger. Messages with less than three valid words were eliminated. The details of this dataset are described in Table 1.

Table 1: Dataset details

Description	Value
Total number of Weibos	252173
Total number of locations	1088
Average length of Weibos	110.32

4.2 Comparison Methods

We compare the proposed model LGLDA with the following other methods.

- **TF-IDF with K-means clustering (or TF-IDF).** Weibos are firstly presented as tf-idf weighted vectors and aggregated by locations. Then, feature vector of each location is summed by documents and clustered by K-means. The center of each cluster represents a topic, and keywords are extracted by feature weight.
- **LDA model with location aggregation (or LDA).** Topic and word distributions are firstly calculated by standard LDA algorithm with all documents, without considering the geo-tagged locations. Then, the topic distribution of each location is calculated as the average over all documents geo-tagged with this location.
- **Local LDA model (or LocalLDA).** This method is similar to those in previous works but a simplified version in our scenario. Compared to LDA described above, the documents are calculated with geo-tagged locations in LocalLDA. In LocalLDA, there is only a local topic distribution, and other settings are kept the same with the LGLDA model.

4.3 Quantitative Measures

In order to make a comparison between different methods, several quantitative measures are used.

Perplexity is used to evaluate the performance of topic modeling. A lower perplexity score indicates better generalization performance of the model.

$$perplexity(D) = exp\{-\frac{\sum_{d\in D} \log p(w_d)}{\sum_{d\in D} N_d}\} \quad (4)$$

where D is the test collection and N_d is document length of document d.

Two entropies: the average topic entropy of each location and the average location entropy of each topic are used to measure the focusness of discovered topics. Each location should have a compact distribution on topics, while each topic should concentrate on a small set of locations.

$$entropy_{topic} = \sum_L (\sum_K p_k^{(l)} \log p_k^{(l)})/L \quad (5)$$

$$entropy_{location} = \sum_K (\sum_L p_l^{(k)} \log p_l^{(k)})/K \quad (6)$$

where $p_k^{(l)}$ and $p_l^{(k)}$ are the estimated probabilities of topic k for location l and location l for topic k, respectively.

KL-divergence is used to measure the average distance of word distributions of all pairs of topics. The larger the average KL-divergence is, the more distinct the topics are.

4.4 Settings and Results

In our experiments, the number of topics was set at K to 20 for all models (including K-means), α_l and α_g to 0.1, β to 0.1, γ_l and γ_g to 0.5 for all LDA-based models empirically and were run for 500 iterations. For LGLDA model, the local-global weight ratio was determined by model's perplexity as shown in Figure 2. As can be seen, with the gradually increasing value of λ from 0.1 to 20, the perplexity descended first and then ascended, and reached minimum at 0.6. Hence, 0.6 is the best value for λ, and was used in our experiments. Performance on the other two metrics also confirmed this selection.

In order to validate the effectiveness of our model to distinguish between local words and noises, the locality score of Weibo defined as the ratio of the average generation probability of words from the local and global topic distribution was calculated in Equation (7). The higher the score is, the more relevant to local context the Weibos are.

$$Locality(d) = \sum_{w\in d} p(z_{e,w}, e_w = l) / \sum_{w\in d} p(z_{e,w}, e_w = g) \quad (7)$$

Three Weibos were sampled from the collection with distinguishable locality score with location tagged 'Waitan' and shown in Table 2. The result is in accordance with expectation. Weibo with the highest score which includes several landmark names and location specific words is highly rele-

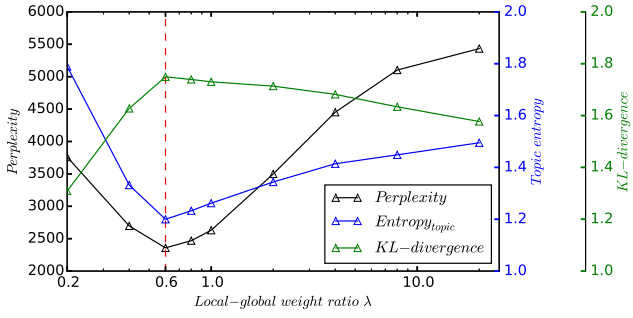


Figure 2: The impact of the value of local-global weight ratio λ on model's performance.

Table 2: Examples of Weibo text with locality score sampled from location tagged Waitan

Weibo text	Locality score
Cruising along Waitan while taking a close-up view of the Oriental Pearl Tower is a worthwhile trip.	7.18
Breakfast at Xitang, lunch at Hangzhou, dinner at Waitan. An incredible day on the run.	0.91
To deal with a difficult customer in the afternoon, I'd have to get up early to do data analysis.	0.02

Table 3: Results of the comparative experiments

Method	Perplexity	Topic entropy	Location entropy	KL-divergence
LDA	6904.11	2.9680	59.5100	**2.2944**
LocalLDA	5679.95	1.5156	31.2292	1.5694
LGLDA	**2357.94**	**1.1998**	**24.2575**	**1.7494**

vant, while the other two containing only few or no location featured words are weakly related or irrelevant.

Table 3 gives the comparative results with other methods. Our model outperforms the others in perplexity, topic and location entropies, and KL-divergence. It achieves much lower perplexity since it can separate local and non-local words, thereby the documents are better organized. With noises filtered out, the discovered topics are more concentrated and representative for local semantics. However, the reason that KL-divergence of our model is lower than the LDA baseline is due to the location relevance constraint.

Table 4: A comparison of TOP3 topics discovered

TF-IDF	LDA	LocalLDA	LGLDA
work, company, mood, women, teacher	love, feeling, mate, inside, teacher	work, mood, love, children, phone	work, company, mood, phone, city
university, teacher, effort, mood, paper	teacher, English, school, exam, culture	university, teacher, library, school, birthday	university, school, library, teacher, student
Waitan, city, night, Oriental Pearl Tower, restaurant	city, Waitan, Shanghai, Oriental Pearl Tower, international	Waitan, hotel, center, Oriental Pearl Tower, financial	Waitan, Oriental Pearl Tower, Chenghuang Temple, Nanjing Road, Huangpu River

Table 4 represents the comparison of the top3 topics discovered. The first topic (2nd row) is composed of broader words which can be viewed as noises, while the other two (3rd and 4th row) contain the semantics of education and tourist attractions. As can be seen, keywords discovered by our LGLDA model achieve the best relevance while keywords by other methods include more or less noises. From our results, Topic1 together with Topic4 accounts for 98.8% of global topic distribution in LGLDA. With noises filtered out, locations are covered by less topics. For example, for Oriental Pearl Tower, the weight of Topic3 in LGLDA is 0.933 compared with a mixed topic constitution discovered by LocalLDA (0.425 for Topic1, 0.414 for Topic3 and 0.161 for others).

5. CONCLUSION

This paper proposes a method to combine local word filtering and geographical topic modeling into one framework. The proposed LDA-based model can effectively distinguish between local words and a variety of noisy daily interests by properly choosing the local-global weight ratio parameter. Results on Weibo collection show the effectiveness of our method over other baselines.

6. ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (Grant No. 61371084 and 91324202). Data and code will be published on Github after this review.

7. REFERENCES

[1] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *the Journal of machine Learning research*, vol. 3, pp. 993–1022, 2003.

[2] H. Li, Z. Li, W.-C. Lee, and D. L. Lee, "A probabilistic topic-based ranking framework for location-sensitive domain information retrieval," in *SIGIR*, pp. 331–338, ACM, 2009.

[3] Z. Yin, L. Cao, J. Han, C. Zhai, and T. Huang, "Geographical topic discovery and comparison," in *WWW*, pp. 247–256, ACM, 2011.

[4] L. Hong, A. Ahmed, S. Gurumurthy, A. J. Smola, and K. Tsioutsiouliklis, "Discovering geographical topics in the twitter stream," in *WWW*, pp. 769–778, ACM, 2012.

[5] Q. Yuan, G. Cong, Z. Ma, A. Sun, and N. M. Thalmann, "Who, where, when and what: discover spatio-temporal topics for twitter users," in *SIGKDD*, pp. 605–613, ACM, 2013.

[6] Y. Liu, M. Ester, B. Hu, and D. W. Cheung, "Spatio-temporal topic models for check-in data," in *ICDM*, pp. 889–894, IEEE, 2015.

[7] L. Backstrom, J. Kleinberg, R. Kumar, and J. Novak, "Spatial variation in search engine queries," in *WWW*, pp. 357–366, ACM, 2008.

[8] Z. Cheng, J. Caverlee, and K. Lee, "You are where you tweet: a content-based approach to geo-locating twitter users," in *CIKM*, pp. 759–768, ACM, 2010.

[9] F. Wu, Z. Li, W.-C. Lee, H. Wang, and Z. Huang, "Semantic annotaion of mobility data using social media," in *WWW*, pp. 1253–1263, ACM, 2015.

[10] https://en.wikipedia.org/wiki/Sina_Weibo.